

associating [a plurality of document] item features with each [said document] item, wherein each of said [document] item features represents one of a plurality of tokens found in said [document] item;

producing a feature vector for each item [said document], wherein each feature vector includes said [plurality of document] item features with a count corresponding to each [document] item feature, said count indicative of the number of times said [document] item feature appears in said [document] item;

multiplying said category vector by said [document] item vector [, in accordance with the mathematical convention of multiplication of a vector by a vector,] to produce a plurality of category scores for each [document] item; and

for each perspective, across multiple perspectives, classifying [a document] an item into a category provided said category score exceeds a predetermined threshold.

2. (Amended) The method of claim 1, wherein [the producing a feature vector step further comprises producing a feature vector for each said document, wherein each feature vector includes said plurality of document features with a count corresponding to each document feature, said count indicative of the number of times said document feature appears in said document, provided said document feature is not a stop word.]

the count includes at least one subcount indicative of the number of times the item feature appears in a particular region of the item, wherein the particular region of the item is a subset of the item taken as a whole; and

for first and second categories, in classifying the item into the second category, substituting the at least one subcount for the count of each item feature, depending on whether the item was classified into the first category.

3. (Amended) The method of claim 1, wherein [the producing a feature vector step further comprises producing a feature vector for each said document, wherein each feature vector includes said plurality of document features with a count corresponding to each document feature, said count

indicative of the number of times said document feature appears in said document, provided said document feature is not FilteredOut.]

for first and second categories, classifying the item into the second category depends at least in part on a determination of whether the item was classified into the first category.

[Independent claim 1a]

4. (Amended) The method of claim 1, wherein the classifying [step] further comprises:
comparing a category score for a first [document] item in a first perspective with a category score for said first [document] item in a second perspective; and
modifying the category score in the first perspective in response to the category score in the second perspective.

[Independent claim 1.a.3]

8. (Amended) The method of claim 4, wherein the comparing [step] further comprises: creating an ordered list of category scores for [each category] all categories in [each perspective] all perspectives;
identifying an item associated with a highest category score in a first perspective;
identifying said item in a second perspective; and
decreasing the category score of said item in said second perspective.

11. (Amended) The method of claim 8, further comprising [the step of] repeating the [locating] identifying and decreasing [steps] for every perspective.

13. (Amended) A method for associating at least one of a plurality of features with at least one of a plurality of categories, said method comprising at least one of manually or automatically associating at least one of said plurality of features to at least a first category, said plurality of features contributing to a decision to classify a document or query item into said at least first category.

14. (Amended) The method of claim 13, further including classifying at least one [document] item into said at least one category, provided the [document] item includes a predetermined number of said plurality of features associated with said category.

15. (Amended) The method of claim 13, further comprising at least one of manually or automatically associating at least one of a plurality of attributes with at least one of said plurality of features, said plurality of attributes contributing to a decision to classify [a document] an item into said at least one category.

19. (Amended) The method of claim 15, further comprising classifying [a document] an item into a category, provided the [document] item does not contain a feature whose association with said category has a RejectConcept attribute

20. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item that contains [a] the feature or a morphological variant of the feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be morphologically variable.

21. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item that contains [a] the feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be morphologically invariant.

22. (Amended) The method of claim 15, further comprising [the step of not] using the presence of a feature in an item for avoiding classifying [a document] the item that contains a morphological variant of a feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be morphologically invariant.

23. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item that contains a feature or a case variant of the feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be case insensitive[. a feature associated with a category which association has an attribute declaring the feature to be case insensitive contributes to a decision to classify a document into said category, provided the document contains the feature or a case variant of the feature].

24. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item that contains a feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be case insensitive.

25. (Amended) The method of claim 15, further comprising [the step of not] using the presence of a feature in an item in avoiding classifying [a document] the item that contains a case variant of a feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be case invariant.

26. (Amended) The method of claim 15, further comprising classifying at least one [document] item into at least one of said categories, provided the [document] item contains a feature whose association with said at least one category has [an] a DirectHit attribute [entitled DirectHit].

27. (Amended) The method of claim 15, further comprising [classifying at least one document into at least one of said categories, provided the document does not contain a feature whose association with said at least one category does not contain the attribute entitled Overlap]. using an Overlap attribute of a first feature to determine whether to use the presence of the first feature in an item, wherein said first feature overlaps a second feature in the item, in classifying the item that contains the first feature into a category.

28. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item containing an overlapping feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be overlap insensitive.

29. (Amended) The method of claim 15, further comprising using the presence of a feature in an item for classifying [a document] the item containing a non-overlapping feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be overlap sensitive.

30. (Amended) The method of claim 15, further comprising [the step of not] using the presence of a feature in an item in avoiding classifying [a document] the item that contains an overlapping feature into a category, provided the feature contains an attribute associated with the category that declares the feature to be overlap sensitive.

31. (Amended) The method of claim 15, further comprising at least one of manually or automatically assigning a weight to the feature, said weight indicative of a degree of association between said [document] item and said category.

32. (Amended) The method of claim 31, further comprising:
determining whether said weight was manually assigned to said feature; and
associating an attribute with said feature that indicates [that] whether the weight was WeightEdited.

38. (Amended) The method of claim 15, further comprising [the steps of]:
determining whether at least one of said plurality of features is a stop word; and
setting an attribute indicating that said feature is a stop word.

39. (Amended) The method of claim 15, further comprising [the steps of]:

at least one of manually or automatically determining a scope of at least one of said plurality of features; and

setting an attribute indicating that said at least one feature is for queries only, or for documents only, or for both.

40. (Amended) The method of claim 15, further comprising [the step of] setting an attribute indicating that said feature is FilteredOut, provided said feature has been manually or automatically filtered out of a classification.

41. (Amended) The method of claim 31, further comprising multiplying said weight by a scaling parameter, provided the decision to classify the [document] item into said category was based on at least one feature automatically associated with the category.

45. (New) The method of claim 1, in which the multiplying includes using at least one attribute of at least one category vector in determining whether to include a document feature of the feature vector in the multiplying.

46. (New) The method of claim 45, in which the using at least one attribute includes using at least one of:

a Stop attribute to indicate whether a feature must constitute something other than a stop word to be included in the multiplying;

a Case attribute to indicate whether a feature must match a letter case specification to be included in the multiplying;

a Stemming attribute to indicate whether a feature includes stemmed word forms to be included in the multiplying; and

a Learned attribute to indicate how a human-specified feature is to be included in the multiplying.

47. (New) The method of claim 45, in which the perspective to which the category relates determines the value of the at least one attribute.

48. (New) In a system including perspectives and categories, each perspective comprising at least one category representative of that perspective, a method for constructing a classifier to classify at least one item across multiple perspectives, the method including:

associating at least one feature with each category, in which each feature is configured for being detected in at least a portion of the at least one item for classification of that item;

determining an initial weight indicating a degree of association between each associated feature and category; and

in which weights for a category are initially related to weights for other categories of the same perspective but are initially substantially unrelated to weights for categories in different perspectives.

49. (New) The method of claim 48, in which the determining the weight indicating a degree of association between each associated feature and category includes using the corresponding feature's distribution in training data items tagged to categories from the same perspective as the category being associated with the corresponding feature.

50. (New) The method of claim 48, in which the determining the weight indicating a degree of association between each associated feature and category includes receiving a user input specifying the weight.

51. (New) The method of claim 50, further including deeming a feature to be unassociated with a category if no user input is received specifying the weight corresponding to the feature.

52. (New) The method of claim 48, further including deeming a feature to be unassociated with

a category if a magnitude of a corresponding weight between the feature and the category does not exceed a predetermined threshold value.

53. (New) The method of claim 52, further including specifying the predetermined threshold value, for each perspective, independent of the predetermined threshold value for other perspectives.

54. (New) The method of claim 48, further including deeming a feature to be unassociated with a category if a number of features associated with the category exceeds a predetermined threshold value.

55. (New) The method of claim 48, further including limiting how many categories an item can be classified into within a particular perspective.

56. (New) In a system comprising perspectives and categories, each perspective comprising at least one category representative of that perspective, the system also comprising weights, each weight indicating a degree of an association between a feature and a category, a method for classifying at least one item across multiple perspectives, the method comprising:

identifying feature instances in the items;

representing, for each item, which features were identified in that item and the number of instances each such feature was identified in that item;

computing, for each item, a category score for each category associated with at least one feature identified in that item, the computing using the weight associating the category and the at least one feature identified in that item;

selecting one or more categories to represent each perspective according to the category scores; and

classifying the items across the selected categories representing the multiple perspectives.

57. (New) The method of claim 56, in which the selecting includes comparing, for each category, the category score to a predetermined threshold value of the perspective represented by the

category, including filtering out those categories with a category score below the predetermined threshold value, and keeping those remaining categories with a category score that equals or exceeds the predetermined threshold value.

58. (New) The method of claim 56, further including, for each perspective, limiting the number of categories for that perspective to be less than or equal to a category count limit for that perspective, and in which the selecting one or more categories to represent each perspective includes eliminating categories in excess of the category count limit for that perspective based on their relatively lower category scores.

59. (New) The method of claim 56, in which the items include representations of documents or queries.

60. (New) The method of claim 56, in which the identifying features includes identifying vocabulary relevant to at least one category.

61. (New) The method of claim 56, in which the identifying feature instances in the items includes identifying features instances in documents, and further including using an attribute designating at least one region of the document to which the identifying feature instances is limited.

62. (New) The method of claim 56, further including receiving user input for determining at least one weight.

63. (New) The method of claim 56, further including statistically determining at least one weight using training data.

64. (New) The method of claim 56, further including providing, for each feature, at least one weight associating that feature with a corresponding category, and basing the at least one weight on at least one of:

automated processing of training data;

user-input data; and

a combination of automated processing of training data and user-input data.

65. (New) The method of claim 64, further including specifying, for each perspective, a degree for combining automated processing of training data and human-input data for weights associated with categories representative of that perspective.

66. (New) The method of claim 56, further including modifying weights initially associating features with one or more categories representing a first perspective based on other weights associating features with one or more other categories representing one or more perspectives different from the first perspective.

67. (New) The method of claim 66, in which the modifying weights includes, if a feature's initial weights indicates that the feature is strongly correlated with at least one category in a first perspective and weakly correlated to the categories of different perspectives, then doing at least one of:

reducing the feature's weights to the categories of the different perspectives; and

increasing the feature's weight to the category of the first perspective.

68. (New) The method of claim 56, further including incorporating a dependence between an item's category score for a category representing a first perspective and the item's category score for one or more other categories representing one or more perspectives different from the first perspective.

69. (New) The method of claim 58, in which the incorporating the dependence includes, if the item's category score for a category representing a first perspective equals or exceeds a threshold value, then inhibiting classification of the item to one or more other categories representing one or more perspectives different from the first perspective.

70. (New) The method of claim 68, in which the incorporating the dependence includes, if the item's category score for a category representing a first perspective equals or exceeds a threshold value, then reducing the item's category score for one or more other categories representing one or more perspectives different from the first perspective.

71. (New) The method of claim 56, in which the items are documents, and in which the classifying the items includes limiting the regions of the document that are used for classifying the documents to at least one category in a second perspective based on a characteristic of a classification of the document to at least one category in the first perspective.

72. (New) In a system for classifying items to categories, a method including:
receiving user-input defining all associations between classification features and categories;
and
statistically determining weights corresponding to the user-defined associations, each weight indicating a degree to which the association's feature identifies the association's category and discriminates against other categories.

73. (New) The method of claim 72, further including identifying candidate features in documents or queries, in which the candidate features include words or phrases in the documents or queries.

74. (New) The method of claim 72, further including providing attributes for the features.

75. (New) The method of claim 74, in which the providing attributes includes providing at least one of:

an Exact Match attribute to indicate whether a match to the feature requires both matching case and matching a stemming form;

a Case attribute to indicate whether a match to the feature requires matching case;

a Stemming attribute to indicate whether a match to the feature requires matching a stemming form; and

an EmbeddedTermsAllowed attribute to indicate whether a match to a first feature precludes a match to one or more other features embedded within the first feature.

76. (New) The method of claim 72, further including providing attributes for the associations of features to categories.

77. (New) The method of claim 76, in which providing attributes includes providing at least one of:

an Edited attribute indicating whether the association has been specified in a Recorded Evidence Edits table;

a WeightEdited attribute indicating whether a weight of the association was specified or edited by a human user;

a Stop attribute indicating whether the feature is a stop word;

a Scope of Feature attribute indicating whether the association of the feature to the category applies to topic spotting of queries only, topic spotting of documents only, or topic spotting of both queries and documents;

a FilteredOut attribute indicating whether the feature should be disregarded during topic

spotting;

a DirectHit attribute that, if asserted, indicates that a document or query including the feature should be tagged to the category specified in the association between feature and category bearing the DirectHit attribute, regardless of what other features are included in the document or query;

a RejectConcept attribute that, if asserted, indicates that a document or query including the feature should not be tagged to the category specified in the association between the feature and category bearing the RejectConcept attribute, regardless of what other features are included in the document or query;

a Case attribute that indicates whether a to indicate whether a match in to the feature, in a document or query, requires matching case of the feature;

a Stemming attribute to indicate whether a match to the feature, in a document or query, requires matching a stemming form; and

an EmbeddedTermsAllowed attribute to indicate whether a match to a first feature, in a document or query, precludes a match to one or more other features embedded within the first feature, in the document or query.

78. (New) The method of claim 72, further comprising receiving user input for overriding at least one of the computed weights indicating a strength of the association of a feature to a category.

79. (New) The method of claim 78, in which the overriding at least one of the computed weights includes increasing the weight's strength of association of a feature and a category.

80. (New) The method of claim 72, in which the statistically determining weights is based at least in part on how often the associations' features are present in a set of training items.

81. (New) The method of claim 72, in which the determining the weight indicating a degree of

association between each associated feature and category includes receiving a user input specifying the weight.

82. (New) The method of claim 81, further including deeming a feature to be unassociated with a category if no user input is received specifying the weight corresponding to the feature.

83. (New) The method of claim 72, further including receiving user input for determining at least one weight.

84. (New) The method of claim 72, further including statistically determining at least one weight using training data.

85. (New) The method of claim 84, in which the training data consists essentially of at least one document associated with a category.

86. (New) The method of claim 84, in which the training data consists essentially of at least one user-specified association between a feature and a category.

87. (New) The method of claim 72, further including providing, for each feature, at least one weight associating that feature with a corresponding category, and basing the at least one weight on at least one of:

automated processing of training data;

user-input data; and

a combination of automated processing of training data and user-input data.

88. (New) The method of claim 87, further including specifying, for each perspective, a degree for combining automated processing of training data and human-input data for weights associated with categories representative of that perspective.

89. (New) In a system for classifying items to categories, a method including:

receiving user-input creating user-defined associations between classification features and categories;

statistically determining machine-defined associations that are capable of being different from the user-defined associations; and

classifying items to the categories using weights corresponding to the user-defined associations and the machine-defined associations.

90. (New) The method of claim 89, further including controlling relative contributions of the weights corresponding to the user-defined associations and the weights corresponding to the machine-defined associations.

91. (New) The method of claim 90, in which the controlling includes obtaining a generally greater relative contribution of the user-defined associations with respect to the machine-defined associations.

92. (New) The method of claim 89, in which the classifying includes classifying the items to categories spanning multiple perspectives.